

M2M Traffic via Random Access Satellite links: Interactions between Transport and MAC Layers

Manlio Bacco, Tomaso De Cola, Giovanni Giambene, Alberto Gotta

Abstract Machine-to-machine services are witnessing an unprecedented diffusion, which is expected to result in an ever-increasing data traffic load. In this context, satellite technology is playing a pivotal role, since it enables a widespread provisioning of machine-to-machine services. In particular, oil industry, maritime communications, as well as remote monitoring are sectors where the use of satellite communications is expected to dramatically explode within the next few years. In the light of this sudden increase of machine-to-machine data transported over satellite, a more thorough understanding of machine-to-machine service implementation over satellite is required, especially focusing on the interaction between transport and MAC layers of the protocol stack. Starting from these observations, this paper thoroughly analyses the interaction between TCP and the Contention Resolution Diversity Slotted Aloha access scheme defined in the DVB-RCS2 standard, assuming the use of an MQTT-like protocol to distribute machine-to-machine services. A novel TCP model is developed and validated through extensive simulation campaigns, which also shed important lights on the design choices enabling the efficient transport of machine-to-machine data via satellite.

I. INTRODUCTION

According to data traffic forecast reports, the volume of data transported by Internet in 2019 will exceed the threshold of 2.0 zettabytes per year, generated by more than trillion of devices. Only a minor portion of the traffic will be generated by PCs, as commonly observed in the recent past; on the contrary, a large quota of Internet traffic will be generated by TVs, tablets, smartphones, and Machine-to-Machine (M2M) devices. In particular, it has been highlighted that M2M traffic will experience a growth rate in the order of 60%. The large amount of traffic contributed by M2M applications will have an important impact on the design of future network architectures and on dimensioning the capacity of the telecommunication infrastructures. In more words, M2M applications are today largely diffused in several terrestrial deployments and have pushed the scientific community to thoroughly investigate network design implications. The survey in [1] underlines that current M2M markets are fragmented, because various vertical M2M solutions have been already designed and implemented: the efforts to identify a common architecture is of a primary concern. Moreover, reference [1] also shows

that typical Internet of Things (IoT)/M2M services exhibit intermittent behaviour (e.g., ON-OFF oscillations), low data-rates, and high traffic burstiness, thus motivating the use of Random Access (RA) schemes. Seminal works about the use of RA for the efficient transport of M2M services via satellite are provided in [2], [3]. The use of satellites is expected to become crucial in the near future [4], [5], because very highly dense networks, such as M2M ones, pose several challenges to terrestrial wireless networks [6]. In fact, the availability of satellite connectivity is of paramount importance in scenarios such as remote plant monitoring, ship tracking service, and aeronautical telemetry transmission, just to mention a few, where a typical terrestrial infrastructure cannot be exploited. From this standpoint, a special note has to be reserved to the case of M2M services distributed via satellite, whose related industry is continuously increasing in size.

The design of efficient RA schemes for satellite communications has captured a great interest from the scientific community and the satellite industry over the last decade [2], [3]. An important example is represented by the second generation of DVB-RCS (DVB-RCS2) [7], specifying the transmission schemes and the protocol architecture for the return link, where Contention Resolution Diversity Slotted Aloha (CRDSA) [8] is proposed as a complementary method to the traditional Demand Assignment Multiple Access (DAMA). CRDSA offers immediate access to satellite capacity without incurring in the reservation delays of DAMA that can penalise the transmission of bursty traffic (e.g., M2M). CRDSA consists in a more sophisticated version of Slotted Aloha (SA), where multiple replicas of the same MAC packet are sent and where accidental collisions are solved by means of Successive Interference Cancellation (SIC). CRDSA relies on the use of two replicas, while the most recent CRDSA++ protocol considers also the use of more than two replicas [9].

The focus of the scientific community has been so far mostly around the implications of RA on physical layer design, only partly addressing its interactions with upper layers (e.g., TCP/IP stack). In particular, the performance analysis of TCP-based higher layer protocols running on top of RA schemes has not yet provided meaningful indications about the boundaries of its applicability, apart from the contribution in [10], which anyway offers a simple and preliminary empirical evaluation. Furthermore, the impact of the collision resolution scheme on TCP congestion control in terms of delay as well as the reaction of TCP in response to TCP segments' collision still needs a rigorous investigation from a theoretical standpoint. This paper attempts to bridge this scientific gap, by analyzing the performance of TCP-based application protocols over CRDSA++, using 3 replicas and running in a DVB-S2/RCS2

(Corresponding author: Manlio Bacco, manlio.bacco@isti.cnr.it)

Manlio Bacco and Giovanni Giambene are with the Department of Information Engineering and Mathematic Science, University of Siena, (e-mail: giambene@unisi.it).

Manlio Bacco and Alberto Gotta are with the Information Science and Technologies Institute (ISTI), National Research Council (CNR), 56124, Pisa, Italy, (e-mail: alberto.gotta@isti.cnr.it)

Tomaso De Cola is with the German Aerospace Center (DLR), Institute of Communications and Navigation, 82234, Oberpfaffenhofen, Germany, (e-mail: tomaso.decola@dlr.de)

network under different traffic load conditions. To this end, a rigorous analytical model of TCP dynamics over RA is drawn, adapting and extending the TCP models in [11]–[13] that cannot inherently apply to the case of RA schemes. The analytical approach will be validated via NS-3-based simulations [14].

The remainder of this paper is structured as follows: Section II surveys the main findings on TCP modeling and interactions with RA schemes, paying special attention to the case of satellite networks. Section III introduces the reference system, while Section IV proposes some refinements that are necessary on long-delay RA links in order to correctly estimate TCP throughput. The performance analysis is carried out in Section VI. Section V offers a simple but effective model to link together burst losses at MAC layer, due to collisions, and segment losses at transport layer. Finally, conclusions are drawn in Section VII.

II. OVERVIEW

A. Related works

In this Section, a survey is presented about the theoretical models available in the literature for estimating the TCP throughput. The literature about TCP modeling can be divided into two main sets, characterized by different packet loss models: bursty losses [11], [15] and independent losses¹ [16], [17]. The TCP behavior is modeled as a *renewal process* for the first time in [11] and then in [17]. A *TCP cycle* is defined as the period of time between to consecutive loss events and it is equivalent to an *epoch* of the corresponding renewal process. The performance of TCP is modeled in terms of *rounds*: an epoch contains m rounds. A round begins with the transmission of W segments, called Congestion Window (CWND), and ends upon the receipt of the Acknowledgments (ACKs) corresponding to these segments. The duration of a round is assumed to be independent of the CWND size and dependent on the Round Trip Time (RTT). The CWND grows at each round, until a loss event occurs; then, TCP enters the recovery phase. Apart from the recovery algorithm that may vary according to the TCP flavor, TCP leaves the recovery phase with or without the expiration of the Timeout (TO), thus entering a new epoch. The first rounds of the new epoch correspond to either a Slow Start (SS) or a Congestion Avoidance (CA) phase, the former if TO has expired, the latter otherwise. In [11], a Triple-Duplicate ACK (TD) is used as loss indication and a TD Period (TDP) (i.e., the period between two TD loss indications) coincides with an epoch of the renewal process. Since our aim is to establish a relationship between the throughput T of a TCP connection and the segment loss probability p , let S_i be the number of TCP segments sent during the i -th epoch and D_i the duration of that period. Thus, we can express the TCP throughput as follows:

$$T = \frac{\mathbb{E}[S_i]}{\mathbb{E}[D_i]}. \quad (1)$$

¹As far as contention-based medium access schemes (e.g., CRDSA) are concerned, losses due to collisions should be considered as independent events: this is empirically proved in Section VI-A2.

The main outcomes in [11], then refined in [18], are the so-called Square Root Formula (SRF) and the Approximated Model (AM). The model in [11] treats segment losses as independent events in different rounds and as correlated events in the same round: each segment delivered after the first segment loss is supposed to be lost as well, independently of its actual outcome. The throughput models in [11] have been derived for TCP Reno recovery algorithm, which can recover a single loss during the Fast Retransmit / Fast Recovery (FR)². However, more recent TCP implementations, like NewReno or Cubic, have more complex error recovery mechanisms that allow recovering multiple losses during the FR phase. Other works have extended SRF and AM in order to account for FR phase' dynamics with Slow But Steady (SBS) flavor [12], [13]. The most prominent innovation in [12] and [13], with respect to [11], is that both bursty and independent losses are modelled.

When specifically targeting the use of TCP in RA schemes for satellites, the work in [19] provides upper and lower bounds of the throughput for a simplified TCP model in a Slotted Aloha (SA) RA channel. Furthermore, the TCP performance over a satellite Diversity Slotted Aloha (DSA) [20] RA medium is analyzed for a single TCP connection in [21]; the authors assume that the MAC layer operates at a target load G^* and that the Network Control Center (NCC) periodically broadcast an *activity factor* (i.e., a throttling parameter) to dynamically adjust the proposed flow control policy applied at Return Channel Satellite Terminals (RCSTs), in order to maximize the TCP throughput. Analogously, [22] analyzes TCP throughput over CRDSA when packet-level Forward Error Correction (FEC) techniques are employed; instead, in this work, we do not use any packet-level FEC techniques. Furthermore, a single TCP connection is modeled through a fluidic model in [22], less precise than the packet level simulator in use in this work. Finally, reference [23] considers a similar scenario, i.e. the use of MQTT in RA satellite channels, but it focuses on a different metric, that is the *completion time*. In fact, the broker is placed close to the subscribers, thus each publisher delivers its data via satellite to the remote broker. Furthermore, a simplified scenario is taken into account: low/medium traffic conditions are under consideration and the impact of packet fragmentation is not analyzed.

B. Main contributions

This work extends the contribution in [24], where the problem was firstly stated and discussed. Here, a complete set of MAC parameters is taken into account in the analysis to better describe the behavior of TCP over a CRDSA-based satellite channel. This work also proposes to use TCP for sending IoT/M2M traffic over CRDSA++, removing the need for any other flow control scheme in use at MAC layer.

To summarize, the main contributions in this work are:

- the analysis of how the TCP congestion control algorithm behaves in a channel dominated by collisions, where TCP

²RFC 6582: The NewReno Modification to TCP's Fast Recovery Algorithm.




Fig. 1: Scenario under investigation: multiple RCSTs acting as rendezvous nodes and delivering data to remote subscribers via RA satellite link

auto-regulates the sending rate to counteract an increase in Segment Loss Rate (SLR);

- a finer model derived from [11]–[13], which accurately matches the TCP steady state throughput over RA satellite channels;
- a simple but effective model establishing a relation between the loss rate at MAC layer and the loss rate at transport layer (cross-layer study);
- finally, considerations on system stability guaranteed by the TCP control congestion algorithm, even at very high loads.

III. SYSTEM ARCHITECTURE AND SETUP PARAMETERS

A. Reference scenario

This work focuses on the applicability of the DVB-RCS2 standard to scenarios in which IoT/M2M traffic is generated from Wireless Sensor Networks (WSNs), similarly to what is considered in [25]. The collected data are *published*, according to the Publish/Subscribe (PUB/SUB) paradigm; a *broker* node collects data and makes them available to *subscribers*. In the scenario depicted in Figure 1, each WSN has a local broker node (also referred to as *cluster head* or *rendezvous* node); the subscribers are connected via a satellite link. Several WSNs are deployed on the ground and each WSN has a potentially large number of sensors, which periodically generate IoT/M2M data. Each broker has no constraints about power consumption and available computational resources; it is connected to a RCST, sending collected data via an RA satellite channel.

The broker is responsible for the delivery of data collected by the WSN. The broker is supposed to have an endless queue of IoT/M2M data waiting for transmission, so that the offered traffic is sustained. Such assumption is valid in several possible application scenarios: for instance, in battle-field communications, where several disjoint teams can use several WSNs sending and receiving data from a remote command center; in maritime scenarios for fleet management; in Flying Ad-Hoc Networks (FANETs), where several resource-constrained Unmanned Aerial Vehicles (UAVs), composing a swarm, can rely on a local rendezvous node or acting as peers to transmit

data to a remote rendezvous node. In these scenarios, we will provide an analytical model in Sections IV and V to predict TCP throughput in supporting IoT/M2M traffic over a satellite RA channel affected by segment losses due to collisions.

According to ETSI (European Telecommunications Standards Institute) and IETF (Internet Engineering Task Force), the Constrained Application Protocol (CoAP)³ is the reference protocol for M2M⁴. Another IoT/M2M protocol, MQTT⁵, developed by IBM in 1999, is the main opponent to CoAP, with a large history of successful deployments. The two protocols have several differences. For instance: (i) MQTT is TCP-based, instead CoAP is UDP-based; (ii) MQTT has (virtually) no limits on the payload size, while CoAP is limited to 1280 bytes; (iii) MQTT is based on a PUB/SUB paradigm, while CoAP provides a request/response interaction model and is REST-oriented (Representational State Transfer). CoAP has been designed with constrained devices in mind, and this explains the lightweight transport protocol in use and the limitation on the payload size, among others.

MQTT has been selected as reference protocol in this work because it is more common in real installations. Furthermore, the PUB/SUB paradigm implemented by MQTT foresees the use of a broker node (as already considered in the architecture in Figure 1), which acts similarly to a Split-TCP PEP (Performance Enhancement Proxy), widely used in satellite networks.

We suppose that no data is lost (reliable delivery) in the local WSN among the IoT/M2M nodes and the local rendezvous node. The subscribers receive data from the rendezvous node via an RA satellite link. An MQTT-like application TCP-based protocol is assumed at the application layer, ensuring a reliable delivery of data from rendezvous nodes to subscribers. The only different assumption is that the underlining TCP connection is open only at the very first data transmission, to avoid a continuous three-way handshake (3WHS) between a rendezvous node and one or more subscribers for the transmission of each data burst. The performance metric under consideration is the throughput at transport layer that corresponds to the throughput experienced by the application itself, apart from the overhead introduced at the application layer. The coexistence of TCP-based and non-TCP-based traffic in the satellite return link is out of the scope of this work and is left for future studies.

B. Satellite system

Let us consider a satellite system and a finite number N of RCSTs, sending data via a satellite RA return link. The RA protocol is CRDSA++ [9]. Time is slotted and each super-frame contains one or more frames; according to DVB-RCS2 terminology, a frame on a return channel is also referred to as RA block. In our configuration, an RA block spans the whole superframe and is composed of n time-slots, or *transmission opportunities*. Each RCST can exploit any transmission opportunity to send data, also referred to as *bursts*, over the

³IETF RFC 7252.

⁴ETSI 102 690 technical specifications indicate CoAP as reference protocol for M2M scenarios.

⁵MQTT v3.1 protocol specification, IBM developerWorks Technical Library, 2010.

return channel. In this work, each transmission opportunity can be used with probability equal to 1, if any data is ready for transmission. At layer 2, the Return Link Encapsulation (RLE) protocol encapsulates the datagrams coming from the network layer. RLE handles *encapsulation*, *fragmentation* and *frame packing* procedures. According to the upper-layer datagrams size, an RLE data unit can carry: a fraction of a datagram, the whole datagram or more than one datagram, according to the respective lengths. Let r denote the length of an RLE data unit measured in IP datagrams that can be packet together, while $f \geq 1$ is defined as the number of IP datagrams that are lost in a single MAC-layer collision event on the RA channel. In case of a collision, $f = \lceil r \rceil$ IP datagrams must be retransmitted. We assume no IP fragmentation in the following, so that one IP datagram carries a whole TCP segment.

We assume that the RCSTs are power balanced⁶. It is worth remarking that power-balancing is the worst condition RCSTs can face when using RA schemes, because the power unbalancing improves the benefits offered by the *capture effect*, thus improving the performance of SIC.

The DVB-RCS2 standard defines two burst lengths (536 and 1616 symbols, respectively) to transport user data that can be transmitted by means of a rich set of waveforms (WFs), so as to operate in different channel conditions and to support various traffic types⁷. In more detail, the redundancy added by channel coding schemes as well as the modulation order can be tuned according to the channel conditions: DVB-RCS2 provides several predefined configurations, or WFs, in order to satisfy reference Bit Error Rate (BER) figures. As a result, the time-slots of an RA block can accommodate different amount of data: a burst is defined as the amount of data fitting a time-slot. This feature can be exploited to configure DVB-RCS2 devices so as to handle different traffic types, such as small-payload messages characterising M2M-like traffic, or conversely larger ones. In the present work, WF 14 and WF 3 configurations are selected (see Table I for more details about these WFs). Apart from the chosen WFs, all the other parameters are equal in the scenario under consideration. The RA block of WF 14 is composed of time-slots, each carrying 188 bytes; the RA block of WF 3 is composed of time-slots, each carrying 38 bytes. Note that WF 3 is one of the most robust and used waveforms, relying on Quadrature Phase Shift-Keying (QPSK) 1/3, while WF 14 uses QPSK 1/2. The choice of these two WFs stems from the fact that the corresponding volume of data being transported can be considered as representative of typical M2M message lengths.

C. TCP segment size and fragmentation at MAC layer

Two different Maximum Segment Sizes (MSSs) are under consideration: 23 and 173 bytes, meant to exactly fit the payload size provided by WF 3 and by WF 14, respectively.

⁶In the reference scenario, RCSTs are uniformly distributed within the same beam. We assume that an uplink power control scheme is implemented, so that the power of the signal received from each RCST is approximately the same.

⁷Note that DVB-RCS2 also defines two additional burst lengths (664 and 262 symbols, respectively) transported by WF 1 and WF 2, which are however used only for control messages. Hence, they are not considered in this work.

TABLE II: Fragmentation at MAC layer

MSS + headers [bytes]	WF ID	$r = \frac{RLE\ payl.(WF)}{MSS+head.}$	$f = \lceil r \rceil$
23 + 15	3	1	1
23 + 15	14	5.7	6
173 + 15	3	0.175	1
173 + 15	14	1	1

IoT/M2M terminals generate low data-rates, thus the MSSs here considered are enough to periodically transmit small amount of data. Since we are dealing with small IoT/M2M payloads, RObust Header Compression (ROHC) is used to reduce the impact of TCP/IP overhead. ROHC is applied prior to transmitting data on the satellite return link. The payload size of the data stream is assumed to be constant, thus the 52 bytes of TCP/IP headers⁸ reduce to 7 bytes for DATA segments and to 6 bytes for ACK segments⁹. The compression ratio for DATA is larger than 85%, thus effectively reducing the overhead.

Several scenarios are possible, depending on the chosen MSS and WF, because of the fragmentation that can take place at MAC layer prior transmitting a TCP segment, as described in Table II, where we consider that 15 bytes of headers (comprised TCP/IP headers plus layer-2 overhead) have to be summed to MSS. Three scenarios are possible, depending on the time-slot payload size:

- 1) the MSS has the same length ($r = 1$): one TCP segment exactly fits one MAC burst, thus a collision event causes a single segment loss, i.e., $f = 1$;
- 2) the MSS has a larger length ($r < 1$): one TCP segment fits into several MAC bursts. A collision event causes a single segment loss, because losing a fraction of the segment is equivalent to lose the whole segment, i.e., $f = 1$;
- 3) the MSS has a smaller length ($r > 1$): several segments fit into a single MAC burst, thus a collision event causes $f = \lceil r \rceil > 1$ segment losses.

The latter case is discussed separately because it violates the hypothesis of uncorrelated losses in Section IV. Moreover, we will show in Section VI that this configuration has a poor performance, so that any configurations leading to $f > 1$ should be avoided, to reduce the possible waste of costly resources.

D. Interactions between TCP and the underlying RA scheme

The traffic arrival process for each RCST is driven by a sustained and continuous load at application layer, because each rendezvous node continuously collects IoT/M2M data from the WSN and sends them to the remote subscribers. The DVB-RCS2 standard specifies a layer-2 load control mechanism (see Section VI-C) aimed at keeping the system running at a target operating point G_T . However, in this work, the actual operating point G is determined by the TCP congestion control algorithm because it already provides a load control mechanism, even if it is meant for avoiding congestion

⁸40 bytes for TCP/IP headers plus 12 bytes for TCP options.

⁹See Section 4.4 of RFC 6846.

TABLE I: Details on the DVB-RCS2 waveforms in use

Waveform ID	Burst length [symbols]	Payload length ([B], [symbols])	Mapping scheme	Code rate
3	536	38, 456	QPSK	1/3
14	1616	188, 1504	QPSK	1/2

events and not for dealing with RA collisions; because of this, the DVB-RCS2 load control algorithm is not adopted. Every time a collision event causes a burst loss, f TCP segments are lost, leading to the retransmission of the segments and to a reduction of the TCP sending rate.

TCP NewReno deals with segment losses relying on: (i) the use of the FR mechanism; (ii) the TO event, triggered if FR is not successful or entering FR is not possible. The CWND is, then, set to a lower value and the Slow Start Threshold (SST) is set accordingly. In both cases, the sending rate is reduced to counteract further loss events. In this work, the satellite return/forward links are supposed operating in clear-sky conditions¹⁰ in order to focus only on segment loss events caused by collisions (ACKs are assumed to be always correctly received). No Automatic Repeat-reQuest (ARQ) algorithm is in use at MAC layer. Each TCP segment sent to the lower layers (i.e., network and MAC layers) is queued into a large finite buffer, so that buffer losses can be neglected w.r.t. collision losses.

This paper analyzes the capability of TCP to self-regulate its transmission rate; given N TCP flows belonging to N distinct RCSTs, the resulting offered MAC load per RCST is given by $\lambda = G/N$, i.e., the TCP flows contribute in achieving the maximum throughput at a load level G that depends on the dynamics of the congestion control algorithm: this aspect will be deepened in Section VI-C.

IV. ANALYTICAL MODEL OF TCP NEWRENO

This Section presents some analytical refinements to the throughput estimation model for TCP NewReno on top of RA. The throughput estimation models anticipated in Section II-A do not accurately match the empirical results of the simulation campaign, as shown in Section VI, thus motivating the refinements resented here. We stress that the throughput estimation model proposed in this Section is meant for TCP over an RA satellite link, and it has not been tested under different conditions. In this analysis, no assumptions are made on the specific RA protocol in use. Furthermore, given the IoT/M2M traffic profile under consideration, different TCP flavors may show basically the same performance level, because of limited CWND size and reduced available bandwidth, the latter being a typical case of satellite RA links.

The new model presented in the following, namely *NewRenoSAT*, is firstly developed in Section IV-A assuming that all loss events are identified by TD ACKs, and then in Section IV-B handling both TD ACKs and TOs (the so-called *full model*).

A. Throughput model if loss indications are TD ACKs only

In this Section, we derive the model for estimating the TCP throughput in absence of TOs, thus only considering Congestion Avoidance (CA) and FR phases.

The throughput T of a TCP flow can be estimated by analyzing a TCP cycle, an epoch of the renewal process. Let S_i^{CA} and S_i^{FR} be the number of segments successfully transmitted during the CA and FR phases of the i -th cycle, respectively, and S_i^{CAFR} the number of segments sent during a whole cycle, here denoted as CAFR (Congestion Avoidance - Fast Retransmit / Fast Recovery):

$$S_i^{CAFR} = S_i^{CA} + S_i^{FR}. \quad (2)$$

Let D_i^{CA} and D_i^{FR} denote the time duration of CA and FR periods and D_i^{CAFR} the duration of the i -th CAFR cycle:

$$D_i^{CAFR} = D_i^{CA} + D_i^{FR}. \quad (3)$$

The average throughput of a flow is given by (1). During the CA phase, the receiver sends one ACK every b segments it receives (*delayed ACK* feature¹¹), causing the CWND to increase linearly with a slope of $1/b$ segments per round, until the first segment loss occurs. Let us denote by α_i the first lost segment in the i -th cycle and by X_i the round where this loss occurs. The total number of segments sent in i -th cycle is $S_i^{CAFR} = \alpha_i + \gamma_i$, where γ_i is the number of segments sent between the first loss α_i and the last one in the drop window W_i . A *drop window* is defined as a CWND where a loss event has occurred. S_i^{CAFR} is defined as in Eq. (2) in [12]:

$$S^{CAFR} = \mathbb{E}[\alpha] + \mathbb{E}[\gamma], \quad (4)$$

where $S^{CAFR} = \mathbb{E}[S_i^{CAFR}]$. The expected number of segments sent in a cycle, having k rounds, up to α_i , is given in [13] as:

$$\mathbb{E}[\alpha] = \sum_{k=1}^{\infty} k(1-p)^{k-1}p = \frac{1}{p}, \quad (5)$$

where p represents the average loss event rate. A *loss event* is said to happen when one or more losses occur in a CWND, thus triggering the entering in FR (or triggering a TO, as in Section IV-B). The first segment loss in a CWND marks the beginning of a *loss event*. The condition for entering the FR phase is to successfully deliver at least three segments in the drop window W_i . Hence, let us assume $\delta_i \geq 1$ losses with rate q over the remaining $(W_i - 3)$ segments, where q corresponds to the average SLR. The number of losses δ_i follows a binomial distribution over a drop window; in

¹⁰DVB-RCS2 offers adaptive MODCOD (Modulation and Coding) to fulfill a target BER requirement. A quasi-error-free link assumption can be made by carefully choosing the target BER.

¹¹As RFC 1122 suggests, the amount of traffic from the receiver to the sender should be reduced by sending a single ACK every b segments.

particular, let us consider the probability distribution of δ_i losses over $(W_i - 3)$ segments, conditioned on $\delta \geq 1$, as:

$$\begin{aligned} \text{Prob}\{\delta = j \mid \delta \geq 1, W_i > 3\} &= B[W_i - 3, j] \\ &= \binom{W_i - 4}{j - 1} (1 - q)^{W_i - 3 - j} q^{j - 1}, j \in [1, W_i - 3]. \end{aligned} \quad (6)$$

Analogously to (5), the average number of segments sent between two consecutive losses is $1/q$. Therefore, if δ_i losses are assumed, the average number of segments among δ_i losses in the same CWND is $(\delta_i - 1)/q$. The expected value of γ can be calculated as follows:

$$\begin{aligned} \mathbb{E}[\gamma] &= \mathbb{E} \left[\sum_{j=1}^{W_i - 3} \frac{(j - 1)}{q} (q B[W_i - 3, j]) \right] \\ &= q (\mathbb{E}[W] - 4), \end{aligned} \quad (7)$$

where $\mathbb{E}[W]$ is the expected value of the CWND size, under the assumption of steady state, and $(q B[W_i - 3, j])$ is the joint probability of the first loss α_i and of the other $(j - 1)$ losses in the drop window. Note that $\mathbb{E}[\gamma]$ is defined only in presence of the losses. If just a single segment loss occurs, then $\mathbb{E}[\gamma] = 0$; otherwise, if two or more segment losses occur, then $\mathbb{E}[\gamma] > 0$. When $q \rightarrow 1$, $\mathbb{E}[\gamma] \rightarrow (\mathbb{E}[W] - 4)$, i.e., the number of segments sent in a drop window after the first three segments allowing for TD ACKs and the first loss α ; conversely, when $q \rightarrow 0$, $\mathbb{E}[\gamma] \rightarrow 0$.

In order to derive the average number of rounds in a cycle, namely $\mathbb{E}[X]$, and the average CWND size, let us consider the evolution of CWND as a function of the number of rounds. According to [13] and [18], we have:

$$W_i = \frac{W_{i-1}}{2} + \frac{X_i}{b} - 1, \quad (8)$$

thus we can write the following relation, at regime, among average values:

$$\mathbb{E}[X] = b \left(\frac{\mathbb{E}[W]}{2} + 1 \right). \quad (9)$$

The expected number of segments sent during the CA phase is determined in [18] as:

$$S^{CA} = \frac{\mathbb{E}[X]}{4} 3 \mathbb{E}[W] + \mathbb{E}[\beta], \quad (10)$$

where β_i is the number of segments sent in the last round of the i -th cycle. $\mathbb{E}[\beta] \approx \mathbb{E}[W]/2$ because we assume β to be uniformly distributed in $[1, W_i - 1]$.

The number of segments $S_{i,k}$ sent at each round k of the i -th FR cycle, considering the Partial Window Deflation¹² mechanism, can be derived from [26] as:

$$S_{i,k} = \max \left(0, W_i/2 - \delta_i + k - 1 \right), 1 \leq k \leq \delta_i. \quad (11)$$

The entering in an FR phase is due to the reception of a TD. When a partial ACK is received, CWND is decreased by the amount of data acknowledged and increased of one segment size. Therefore $S_{i,k+1} = S_{i,k} + 1$. $S_{i,k}$ has a maximum value for $k = \delta_i$, then $\max(S_{i,k}) = S_{i,\delta_i} = W_i/2 - 1$. Thus, the

expected total number of segments sent in an FR cycle comes from (11) as follows:

$$\begin{aligned} S^{FR} &= \mathbb{E} \left[\sum_{k=1}^{\delta_i} S_{i,k} \right] \\ &= \begin{cases} \frac{1}{2} (\mathbb{E}[\delta] \mathbb{E}[W] - \mathbb{E}[\delta] - \mathbb{E}[\delta]^2), & \text{if } \mathbb{E}[\delta] < \mathbb{E}[W] \\ 0 & \text{if } \mathbb{E}[\delta] = \mathbb{E}[W]. \end{cases} \end{aligned} \quad (12)$$

In the case of a single loss, i.e., $\delta_i = 1$, $S^{FR} = \mathbb{E}[W]/2 - 1$. After the first loss, the remaining $(\delta_i - 1)$ segments are lost over $(W_i - 4)$ segments. Hence, we can write the following formula to characterize $\mathbb{E}[\delta]$, the expected value of the number of losses δ_i :

$$\mathbb{E}[\delta] = \mathbb{E} \left[\sum_{j=1}^{W_i - 3} j B[W_i - 3, j] \right] \approx 1 + (\mathbb{E}[W] - 4) q. \quad (13)$$

Assuming that collisions on an RA channel are independent events and $f = 1$, then the segment losses are independent events, too. Instead, if $f > 1$, each collision in the channel causes a burst of segment losses of average length in the interval $[1, f]$, thus TCP segment losses cannot be considered as independent ones. Further details on this are in Section VI-A2.

Using (10) and (12), we can express the expected number of segments sent during a CAFR cycle, i.e., the expected value of (2), as follows:

$$S^{CAFR} = S^{CA} + S^{FR}. \quad (14)$$

Then, the expected value of CWND, $\mathbb{E}[W]$, can be obtained by equating (4) and (14). In particular, we obtain a second order equation in $\mathbb{E}[W]$, whose only positive solution can be expressed as:

$$\mathbb{E}[W] \approx \Phi + \sqrt{\Phi^2 - \frac{60pq - 8p - 8}{p(8q + 3b)}}, \quad (15)$$

$$\text{where } \Phi = \frac{22q - 3b - 4}{8q + 3b}.$$

Equation (15) explicitly takes b into account, differently from Eq. (14) in [12], where $b = 1$ is assumed. We will show in Section VI that, because of the finer estimations of S^{FR} and $\mathbb{E}[\delta]$, the expected value of CWND in (15) is closer to the simulated one than the other theoretical models under consideration [11]–[13].

Each RTT is supposed as a r. v., whose value does not depend on CWND, as in [11]–[13]. We recall that:

$$D^{CA} = \mathbb{E}[X] + \mathbb{E}[D^\beta], \quad (16)$$

where D^{CA} is the average duration, in number of rounds, of the CA phase, including the time spent waiting for entering in FR. Moreover, D^{FR} is the average duration, in number of rounds, of the FR phase:

$$D^{FR} = \mathbb{E}[\delta]. \quad (17)$$

In fact, NewReno recovers one lost segment per round. Finally, $\mathbb{E}[D^\beta]$ is the expected time, in number of rounds, after the

¹²RFC 6582: The NewReno modification to TCPs fast recovery algorithm.




Fig. 2: Evolution of CWND during CA and FR phases

first loss and before entering FR; it can be approximated as $\mathbb{E}[D^\beta] \approx 1/2$. From (16) and (17), the expected duration of a cycle $D^{CAFR} = \mathbb{E}[D_i^{CAFR}]$ [s], results as:

$$D^{CAFR} = \mathbb{E}[RTT](D^{CA} + D^{FR}). \quad (18)$$

Figure 2 shows the evolution of CWND during CA and FR phases and the duration of the aforementioned phases. A CAFR cycle ends with the end of the FR phase because of one or more segment losses. The duration of the last part of the CA phase is D^β RTTs: TD ACKs are received, prematurely ending the last round of the CA phase. The two segments lost in Figure 2 are recovered during the FR phase immediately following, which lasts $D^{FR} = 2$ RTTs. The CWND size is $(W_i/2 + 3)$ at the start of the FR phase and $W_i/2$ at the end of the FR phase.

Combining (4) and (18), we can finally write the throughput expression [segments/s], neglecting TOs, as:

$$T_{noTO} = \frac{(\mathbb{E}[W] - 4)q + \frac{1}{p}}{\mathbb{E}[RTT] \left\{ b\left(\frac{\mathbb{E}[W]}{2} + 1\right) + \frac{1}{2} + [1 + (\mathbb{E}[W] - 4)q] \right\}} \quad (19)$$

where $\mathbb{E}[W]$ is expressed in (15).

B. Full Model

This section extends the model in Section IV-A to include also TOs as loss indications. Each cycle consists of either a CA or a CAFR period followed by a TO with an SS period. TCP NewReno may experience a TO either during CA or during FR, given that a loss event (L) occurred with probability p . The former transition occurs with probability $p_{TOCA} = P\{TO|CA, L\}$, when TCP does not receive enough duplicate ACKs to trigger FR, while the latter transition occurs with probability $p_{TOFR} = P\{TO|FR, L\}$, when retransmitted segments are lost during the FR phase. TCP can experience a TO in the CA phase when more than $(W - 3)$ segments are lost in a drop window. The conditional probability p_{TOCA} is given by:

$$p_{TOCA} = \mathbb{E} \left[\sum_{j=W_i-2}^{W_i} \binom{W_i-1}{j-1} (1-q)^{W_i-j} q^{j-1} \right]. \quad (20)$$

A TO occurs during FR if any of the retransmitted segments is lost. This condition can be approximated by assuming that if a loss event occurs in FR, then also the retransmitted segment is lost, thus triggering a TO [12]. For δ losses in the drop window, NewReno takes δ RTTs to recover the lost segments, during which it sends S^{FR} segments. The probability that the i -th segment is lost, given that the previous $(i-1)$ segments are successfully delivered, is $(1-p)^{i-1}p$. Therefore, according to [12], it follows that:

$$p_{TOFR} = \mathbb{E} \left[\sum_{j=1}^{W_i-3} B[W_i-3, j] \left[1 - (1-p)^{j \frac{W_i}{2}} \right] \right]. \quad (21)$$

During a TO, TCP does not send any segments. According to [11], the duration of a TO period is given by:

$$D^{TO} = RTO \frac{1 + \sum_{j=0}^5 2^j p^{j+1}}{1-p}, \quad (22)$$

where the initial value of the RTO parameter of TCP is 2 or 3 [s], in recent implementations¹³.

In the SS phase, the initial CWND size is one and it grows until the SST, $W/2$, is reached. Therefore, in the last round of SS, TCP transmits $\mathbb{E}[W]/2$ segments, on average, then enters the CA phase. The rate of increase of the CWND, when the *delayed ACKs* feature is in use, is $x = (1 + \frac{1}{b})$ on an RTT basis [27]:

$$\begin{aligned} S^{SS} &= \mathbb{E} \left[x^0 + x^1 + \dots + x^{\log_x \frac{W_i}{4}} \right] = \mathbb{E} \left[b \left(x \frac{W_i}{4} - 1 \right) \right] \\ &= b \left(x \frac{\mathbb{E}[W]}{4} - 1 \right). \end{aligned} \quad (23)$$

S^{SS} represents the expected value of the number of segments sent during the SS phase as in [12], where we consider the segments of the last round of the SS phase as being part of the CA phase. D^{SS} represents the duration of the SS phase as proposed in [12], but here explicitly considering b :

$$\begin{aligned} D^{SS} &= \mathbb{E} \left[RTT_i \left(\log_x \frac{W_i}{4} + 1 \right) \right] \\ &\approx \mathbb{E}[RTT] \left(\log_x \frac{\mathbb{E}[W]}{4} + 1 \right). \end{aligned} \quad (24)$$

Note that the expected value in (24) is applied to a non-linear expression. We have already assumed RTT_i and W_i as independent r. v.. Moreover, if the distribution of W is concentrated around its mean, as shown in Section VI, then the approximation on the derivation of the expected value in (24) is acceptable.

The final expression for the throughput estimation, when considering also TOs, is given in (25). It is worth remarking that some parameters, for instance $\mathbb{E}[W]$ provided in Section IV-A, are used also here, but they only consider CA and FR phases, without taking SS into account, thus providing an approximated estimate. The formulation in (25) is presented in its

¹³As in RFC 6298, Appendix A.

$$T_{Full} = \frac{(1 - p_{TOFR} - p_{TOCA})(S^{CA} + S^{FR}) + p_{TOCA}(S^{SS} + S^{CA}) + p_{TOFR}(S^{SS} + S^{CA} + S^{FR})}{(1 - p_{TOFR} - p_{TOCA})(D^{CA} + D^{FR}) + p_{TOCA}(D^{SS} + D^{CA} + D^{TO}) + p_{TOFR}(D^{SS} + D^{CA} + D^{FR} + D^{TO})}. \quad (25)$$

whole expression, without any approximations. We will show in Section VI that (19) represents a throughput estimation that shows a very small error, so that the improvement provided by (25) is negligible.

V. THE RELATION BETWEEN BURST LOSS RATE AND EVENT LOSS RATE: THE BLR MODEL

In this Section, a simple but effective model is proposed in order to glue together the Burst Loss Rate (BLR) value experienced at MAC layer and the p and q values at transport layer. BLR is the rate of losses due to collisions. In [8], an analytical recursive model is proposed to derive an upper bound for the throughput of 2-CRDSA, but, to the best of the authors' knowledge, no analytical model is available for estimating throughput and BLR of CRDSA++.

In the following, we use the term *collision* to refer to MAC bursts marked as lost after SIC, thus not recoverable in any way at the receiver. TCP segments are retransmitted if they are lost because of collisions. For the sake of simplicity, in the following we assume that one segment fits exactly into one time-slot (case $r = 1$). We recall that buffer overflow effects are negligible because of the large buffer in use. In the simulation results presented in Section VI, each TCP retransmission corresponds to a collision in the channel. In our study and simulations, we neglect spurious TCP retransmissions.

In what follows, a very simple expression is provided, denoted as the *BLR model*, which links together q , p and BLR . This expression holds if no spurious retransmissions occur and $f = 1$. We recall that p is the rate of loss events, q corresponds to SLR and $\mathbb{E}[\delta]$ is the expected number of segments lost in a loss event. Each collision is responsible of a single segment loss, triggering a retransmission at transport layer. Thus, we can write:

$$q \equiv BLR \quad (26)$$

and, for a finite CWND size:

$$p = \frac{q}{\mathbb{E}[\delta]} \approx \frac{q}{1 + (\mathbb{E}[W] - 4)q} \approx \frac{BLR}{1 + (\mathbb{E}[W] - 4)BLR}. \quad (27)$$

Plugging (26) and (27) in (19) and (25), we obtain the throughput estimation as a function of BLR. Following this approach, it is sufficient to know the collision rate at MAC layer, i.e., BLR, to obtain q and p rates to derive a TCP throughput estimation. In Section VI-B, the *BLR model* is validated against simulation results.

VI. PERFORMANCE EVALUATION

The performance evaluation shown here is based on the NS-3 simulator, implementing a DVB-S2/RCS2 satellite network by means of the modules described in [28]. The TCP flavor in use is NewReno. The available implementation in NS-3 is based on the SBS variant. Note that, as shown in [13], the performance of SBS is not significantly different from that of

TABLE III: System settings

Name	Value
TCP flavor	NewReno
TCP MSSs	23 / 173 [B]
TCP/IP headers size	7 [B] (w/ ROHC)
Initial RTO value	2 [s]
RA scheme	CRDSA++ (3 replicas)
RA blocks per superframe	1
RA block duration	13 [ms]
Time-slots per RA block, WF 14	64
Time-slots per RA block, WF 3	194
Bandwidth	8012820 [Hz]
Roll off	0.2
Carrier spacing	0.3 [Hz]
Nominal Round Trip Time	0.52 [s]

the Impatient variant in the conditions here under consideration, i.e., limited congestion window sizes and independent loss events. The main settings of the simulator are provided in Table III. The results here intend to assess how changing WF and MSS values can alter the overall performance. The main aim of this study is in clarifying if TCP is a suitable choice as transport protocol over a long delay RA link and, if this is the case, which choice of parameters leads to the highest throughput.

We recall that the G^* load, i.e., the normalized MAC load level achieving the highest throughput at MAC level, depends on the number of available time-slots per RA block [29], among other factors. Thus, while keeping the other parameters unmodified, the choice of a different waveform can change the number of available time-slots per RA block, shifting G^* . For instance, when 64 time-slots are available, the G^* value of CRDSA++ is ≈ 0.7 [bursts/time-slot]; when 194 time-slots are available, the G^* value of CRDSA++ is ≈ 0.78 [bursts/time-slot] [29]. The overall system has been tested at several increasing load levels, before, close to and after G^* . Two crucial aspects are characterized in this work: the MAC operating point \hat{G}_T when NewReno is in use over an RA satellite channel, and the system stability; both issues are discussed in Section VI-C. To the best of the authors' knowledge, the characterization of \hat{G}_T is not available in literature for CRDSA++, if we exclude our preliminary contribution in [24].

The focus in what follows is on normalized load levels $G \geq 0.45$, because, up to this point, CRDSA++ throughput is almost equal to the offered load (i.e., the operating point is in the linear part) [9]; beyond this load level, the system works in the quasi-linear part up to the G^* point, and then the MAC throughput dramatically decreases because of a too large BLR. If a DVB-RCS2 satellite system with ARQ enabled (typical setting) works at $G > G^*$ and no load control algorithm is used, the system may experience instability [30].

A. Throughput validation

The NewRenoSAT model described in Section IV is here validated against simulation results and compared with other

approaches in literature (see Section II-A): (i) the contribution in [11], then amended in [18] (here denoted as PFTK model); (ii) the contribution in [12] (here denoted as PWM model); (iii) the contribution in [13] (here denoted as DAKH model). In the following, the provided throughput estimations are obtained by using simulated p and q values because there is no analytical model in literature able to characterize the layer-2 losses when CRDSA++ is in use.

1) Accuracy of the CWND estimation

A critical value to be computed, when estimating the TCP throughput, is the CWND size. While just the expected value is used for the throughput estimation, the CWND exhibits a range of different values during the whole TCP connection lifetime. Figure 3a shows the distribution of CWND obtained from simulations, when WF 14 and MSS = 173 bytes are selected; in the following, this scenario is used as exemplary case. As the load increases (i.e., when the number of RCSTs increases), $\mathbb{E}[W]$ decreases, while the distribution keeps the same shape. A larger G value means higher load on the RA channel and then higher BLRs. Each time two or more RCSTs are involved in a collision, their CWND size is shrunk to a lower value, as in Figure 2 after the two lost segments. Therefore, $\mathbb{E}[W]$ decreases as G increases. In Table IV, the simulation results are compared with the NewRenoSAT model estimations, showing a close match.

The average value of the CWND has a direct impact on the number of segments sent in a CAFR phase, as visible in (14). The distribution of the number of segment sent per CAFR phase is shown in Figure 3b: this is a geometric distribution, as pointed out in [11]. In a CAFR cycle, the sending rate varies if the connection is in the CA phase or in the FR phase: to underline this, Figure 3c shows the distribution of the number of segments sent in FR only. The distribution in the FR phase (in Fig. 3c) is different from the one in the CAFR phase (in Fig. 3b): the CA phase completely dominates the dynamics of a CAFR cycle, and it exhibits an approximate geometric distribution, while the FR phase has quite a negligible impact on the overall distribution, thus the approximation of a geometric distribution still stands. Table IV also shows how the NewRenoSAT model is able to accurately estimate the average number of segments in CA and FR phases, along with their average durations.

2) Average number of segment losses per loss event

A central part of this work is the characterization of the loss process experienced on the RA channel when CRDSA++ is in use. Equation (13) provides the analytical expression of $\mathbb{E}[\delta]$, the expected number of losses per loss event. It ranges from 1 to 1.2 in our simulation results, as shown in Table V, if we exclude the case of WF 14 and MSS = 23 bytes, later discussed. It means that, on average, 1 to 1.2 segments are lost per drop window: the same number of segments is retransmitted in the FR phase, which follows the loss event. Because of the reduced available bandwidth per RCST, the average CWND size is small and, on average, a single loss is experienced per drop window if q is small.

Instead, a strong correlation effect is present in the scenario WF 14 and MSS = 23 bytes: a single time-slot carries up to $r = 5.7$ segments, thus a single collision causes a burst of

$f = 6$ losses at transport layer. Note that r is the maximum number of whole or partial TCP segments into a single time-slot, and not the actual one, because the RLE protocol always tries to ensure the largest occupation of the time-slot, but it is limited by the number of data units actually pending for transmission. It means that, if less than f data units are in the queue, a smaller number will be transmitted in a single time-slot, thus sub-utilizing the available resources. Moving to the performance evaluation, in this case a single collision represents a very stressful event for a connection, because of the burst of losses experienced at TCP layer. This scenario exhibits very low performance because of this effect, when compared with the scenario WF 14 and MSS = 173 bytes and, generally, its performance are worse than any other configuration with $f = 1$. Because of this, we are not intended to the study of the TCP behavior under the hypothesis of bursty losses due to encapsulation effects (i.e., $f > 1$). For the same reason, Table V does not show the $\mathbb{E}[\delta]$ value, whose analytical expression is valid only under the hypothesis of independent losses.

3) Comparison among different approaches for estimating the TCP throughput

In Figures 4, 5 and 6, we compare the aforementioned TCP models for the analysis of TCP throughput over a RA satellite link using the CRDSA++ protocol. The throughput estimation relies on the use of p , q and $\mathbb{E}[RTT]$ values from simulations, as in Table V.


The relative error $\eta = |1 - \frac{T_{est}}{T_{sim}}|$ of the estimations is shown, where T_{est} is the throughput resulting by the analytical models under consideration [11]–[13] and NewRenoSAT, and T_{sim} is the throughput resulting from the simulations. The results here shown are obtained with extensive Montecarlo simulation runs: more than 1 million sent segments per scenario under consideration, with a narrow confidence interval. η is computed for every combination of WF and MSS in use, except for the scenario WF 14 and MSS = 23 bytes ($f > 1$), as explained before. Let us recall that $b = 2$ (delayed ACKs) is used in the simulations in this work; the PMW model does not account for it, so the relative error of this analytical model is larger when compared with the others under consideration.

Figure 4 shows the value of η for the scenario WF 3 and MSS = 23 bytes: 194 time-slots are available for transmission and an RCST can use one time-slot per RA block. A segment fits exactly into one MAC burst ($r = 1$). The NewRenoSAT model proposed in this work show the lowest η value (i.e., the best results), if we except the case with the largest number of RCSTs (i.e., the largest load); the DAKH model and the PFTK one are quite accurate, too, but the latter has a η larger than 0.15 for higher loads and the former tends to overestimate the throughput at medium loads.


In Figure 5 the relative error for the scenario WF 3 and MSS = 173 bytes is shown. In this scenario, a time-slot carries a fraction of the segment: in fact, $r = 0.175$, so that 5.7 time-slots are needed to carry a single segment, which means that the sixth time-slot carries the last fraction of a segment and the initial fraction of the next one. A collision involving the sixth time-slot causes the loss of two TCP segments, thus triggering

TABLE IV: NewRenoSAT estimated values vs. simulation results
WF 14, MSS = 173 bytes, 64 time-slots per RA block


#RCSTs	$\mathbb{E}[W]$		S^{CA} [segments], D^{CA} [s]		S^{FR} [segments], D^{FR} [s]	
	simulated	model	simulated	model	simulated	model
30	40.15	41.8	1367.6, 25.5	1394, 28.4	19.3, 0.62	20.5, 0.66
40	33.4	34.8	1003.7, 21.4	978, 22.4	15.6, 0.58	17.1, 0.63
50	28.5	29.6	721.4, 17.9	716.6, 18.6	13.2, 0.58	14.4, 0.61
60	26.1	26	552.1, 15.6	559.2, 16.5	11.7, 0.58	12.6, 0.61
70	21.6	21.9	433.3, 13.9	406.1, 14.2	10.4, 0.58	10.6, 0.62



(a) Probability Mass Function of the CWND size



(b) Probability Mass Function of the number of segments sent per CAFR cycle



(c) Probability Mass Function of the number of segments sent per FR phase

Fig. 3: Simulation results of the scenario WF 14, MSS = 173 bytes, 64 time-slots per RA block




Fig. 4: Comparison (relative error) between simulations and theoretical approaches for TCP NewReno throughput over CRDSA++ for WF 3, MSS = 23 bytes and 194 time-slots per RA block

the simulation results show a good match with NewRenoSAT and DAKH analytical models anyway.




Fig. 5: Comparison (relative error) between simulations and theoretical approaches for TCP NewReno throughput over CRDSA++ for WF 3, MSS = 173 bytes and 194 time-slots per RA block

two retransmissions for a single collision event. This happens with probability equal to 1/6: because of this, the error on the estimation is larger than the one in Figures 4 and 6, for each model. In this scenario, there is some correlation among losses, so the hypothesis of complete independence among segment losses is not verified. Anyway, the correlation is limited, so

Figure 6 provides the throughput estimation for the scenario WF 14 and MSS = 173 bytes: one segment fits exactly into a single MAC burst ($r = 1$). NewRenoSAT and DAKH models

achieve the best results, even if DAKH shows an average relative error that is larger at high loads. The PFTK model exhibits a larger error, instead; in fact, it does not take the FR mechanism into account.




Fig. 6: Comparison (relative error) between simulations and theoretical approaches for TCP NewReno throughput over CRDSA++ for WF14, MSS = 173 bytes and 64 time-slots per RA block

4) Retransmission TimeOuts

In this Section, the measure ξ is introduced to quantify the rate of TOs in the simulations:

$$\xi = \frac{\mathbb{E}[\#TOs]}{T_s/\mathbb{E}[RTT]}. \quad (28)$$

$\mathbb{E}[\#TOs]$ is the average number of TOs per connection and T_s is the simulation length in seconds, where $T_s/\mathbb{E}[RTT]$ is the number of rounds in each simulation. Therefore, ξ represents the probability that a TO occurs in a round. The numerical values of ξ per scenario under consideration are shown in the last column of Table V.

We observe two effects worth of mentioning: if $r = 1$, ξ has a substantially lower value with respect to the case where $r \neq 1$; then, MAC fragmentation and packing significantly impact on TCP performance. The second effect here discussed is related to the load level: the lowest simulated load level exhibits a ξ value larger than that at higher ones, which is unexpected, at a first glance. At high loads, the collision rate is responsible of sustained segment losses, thus the ξ value reflects this. At medium loads, the collision rate is responsible of moderate losses, thus ξ shows a lower value. In fact, the FR mechanism is particularly efficient in presence of a moderate loss rate, and time-consuming TOs are avoided. Instead, at low loads, another effect is present, which makes ξ larger than expected: the CWND overshooting problem. A collision is a very rare event at low loads, therefore the overshooting problem is possible: a CWND that increases

slowly but continuously over time and whose value goes over the Bandwidth-Delay product increases the probability of spurious TOs.

B. Validating the BLR model

The *BLR model* discussed in Section V is now validated against simulation results. The throughput estimation here presented relies on the use of *BLR* and $\mathbb{E}[RTT]$ values taken from simulations, as in Table V. The scenario WF 14 and MSS = 173 bytes has been chosen because of the very low rate of spurious retransmissions.

Table VI shows the relative error η when comparing the simulation results with the estimations provided by (19) and (25) in the first two columns. Further to this, the accuracy of the *BLR model* is visible in the last two columns, i.e., when plugging (26) and (27) in (19) and (25). The *BLR model* provides good precision: in fact, the larger estimation error is $\leq 7\%$. We recall that the *BLR model* presented in Section V neglects spurious retransmissions, thus causing a small error in the throughput estimation. A consideration is here in order to motivate the need of the *BLR model*: while it can be practical to estimate BLR in real systems, the same cannot be said for p and q when using TCP. Thus, this model offers a simple and elegant way, in the authors' opinion, for estimating the TCP throughput over an RA satellite link.

C. Stability

A critical issue to be taken into account, when dealing with RA protocols, is system stability. CRDSA++, as other RA protocols, has an optimal working point, namely G^* , which exhibits the maximum throughput offered by the MAC protocol, namely T^* . If the system is forced to work at loads $G > G^*$, instability may occur and proper countermeasures are needed, as analysed in [30].

DVB-RCS2 standard encompasses a normative load control algorithm, aimed at keeping the system at a target operating point G_T , where $G_T \leq G^*$. This load control algorithm is not used in our simulator because the work in [31] shows that the aforementioned algorithm needs to be tuned in accordance with each scenario under consideration, and this can be not trivial at all; furthermore, a more complex algorithm would be necessary to target realistic Quality of Service (QoS) requirements. On the contrary, the use of TCP ensures the stability at every load level under consideration, as discussed later in this part, and does not require any tuning. The DVB-RCS2 load control algorithm can limit the magnitude of the load oscillations around G_T according to two different strategies [31]: by reducing the number N of RCSTs or by shifting G_T to a lower value. These oscillations can be responsible for pushing the instantaneous offered load G beyond G^* , thus leading the system towards the instability region; hence, controlling these oscillations is of primary importance. Both strategies offer disadvantages: in an IoT/M2M scenario, a large population of terminals (RCSTs) is common, so that reducing N could be unfeasible. Instead, shifting G_T towards a lower value can be a better solution in our scenario, but the use of a centralized NCC, which periodically monitors the aggregate

TABLE V: Simulation results: the scenarios are described in the first column.
 $\mathbb{E}[RTT]$ is measured in seconds, $Thr.$ is the average simulated throughput [kbps] per TCP connection

WF, MSS	$\#RCSTs$	BLR	r	f	q	p	$\mathbb{E}[\delta]$	$\mathbb{E}[RTT]$	$Thr.$	ξ
3, 23	110	1.1E-4	1	1	1.44E-3	4.5E-4	1.07	0.69	12.50	5.5E-4
3, 23	130	2E-4	1	1	1.70E-3	4.86E-4	1.08	0.68	11.90	3E-5
3, 23	150	5.5E-4	1	1	2.80E-3	6.68E-4	1.11	0.62	9.60	6.2E-5
3, 23	170	9.2E-4	1	1	5.13E-3	9.86E-4	1.16	0.61	7.84	5E-4
3, 23	190	1.1E-3	1	1	7.94E-3	1.20E-3	1.20	0.60	6.86	2.5E-3
3, 23	210	2.9E-3	1	1	1.1E-2	6.3E-3	1.20	0.59	2.90	4E-3
3, 173	110	1.2E-4	0.175	1	4.1E-3	1.5E-3	1.09	2.40	11.94	1E-1
3, 173	130	3.5E-4	0.175	1	8.66E-3	3.5E-3	1.14	1.50	11.90	9E-2
3, 173	150	1.4E-3	0.175	1	2.35E-2	1.1E-2	1.16	0.92	10	7E-2
3, 173	170	2.8E-3	0.175	1	2.94E-2	1.72E-2	1.17	0.77	8.50	5E-2
3, 173	190	4.4E-3	0.175	1	3.3E-2	2.7E-2	1.18	0.72	7.30	6E-2
3, 173	210	6.16E-3	0.175	1	3.5E-2	3.5E-2	1.20	0.69	6.37	7E-2
14, 23	30	2.27E-3	5.7	6	1.3E-2	3.7E-4	-	1.28	10.5	5E-2
14, 23	40	2E-3	5.7	6	1.38E-2	3.8E-4	-	1.275	10.48	4E-2
14, 23	50	2.32E-3	5.7	6	1.64E-2	4.2E-4	-	1.23	10.1	4E-2
14, 23	60	2E-3	5.7	6	1.71E-2	4.3E-4	-	1.19	9.93	4E-2
14, 23	70	1.9E-3	5.7	6	1.76E-2	4.7E-4	-	1.15	9.72	4E-2
14, 173	30	6.41E-4	1	1	8.52E-4	7.45E-4	1.03	0.64	65.75	5.3E-4
14, 173	40	9.26E-4	1	1	1.34E-3	1E-3	1.04	0.60	58.50	1E-4
14, 173	50	1.31E-3	1	1	1.75E-3	1.37E-3	1.04	0.58	51.51	1.4E-4
14, 173	60	1.76E-3	1	1	2.30E-3	1.79E-3	1.05	0.58	45.20	3E-4
14, 173	70	2.3E-3	1	1	3.5E-3	2.28E-3	1.07	0.58	38.11	4.6E-4

load, is needed. The NCC is responsible for choosing the G_T value according to the instantaneous G . Figure 7a, later described, shows that the use of TCP removes the need of a centralized control: TCP can control G_T . Hence, TCP ensures the stability, without requiring any layer-2 control mechanisms that can improperly interact with the TCP congestion control algorithm.

Figure 7a shows the number of RCSTs versus the normalized MAC throughput and normalized MAC offered load, when using TCP on top of WFs 14 and 3. We recall that WF 3 offers a greater number of time-slots carrying small payloads, while WF 14 offers a lower number of time-slots carrying larger payloads. Three load intervals can be read in Figure 7a, w.r.t. the number of RCSTs: the first one, at low loads, up to ≈ 100 RCSTs, where WF 14 offers a larger normalized throughput; the second one, at medium loads, from ≈ 100 to ≈ 350 RCSTs, where WF 3 outperforms WF 14 thanks to the greater number of time-slots; finally, the third interval, at larger loads, from ≈ 350 RCSTs on, where the throughput offered by the two WFs is almost comparable. A consideration is here in order: when using WF 3, TCP behaves in a greedy way, showing a clear load peak when ≈ 150 RCSTs are present, then the load decreases as the number of RCSTs increases. On the contrary, WF 14 shows a more balanced behavior, and a peak is not recognizable. The key difference is the number of available time-slots, which in turns determines the number of active RCSTs per RA block. The number of active RCSTs per RA block determines the SLR that each RCST experiences, because, for a given load G , the corresponding BLR can be estimated, and the BLR model in Section V can be used to relate BLR and SLR.


Let us call \hat{G}_T the working point for TCP NewReno on top of CRDSA++ protocol. According to the y-axis values in Figure 7a, we can see that the following ranges are possible for

the working point: (i) $\hat{G}_T \in [0.45, 0.55]$ for TCP over WF 14 (with 64 available time-slots); (ii) $\hat{G}_T \in [0.58, 0.63]$ for TCP over WF 3 (with 194 available time-slots). When CRDSA++ (3 replicas) is in use and assuming no flow control algorithms, (i.e., TCP is not in use, as well), $G^* \approx 0.7$, when 64 time-slots are available per RA block [29], and $G^* \approx 0.78$, when 194 time-slots are available per RA block [29]. Thus, as Figure 7a shows, $\hat{G}_T \leq G^*$: the segment loss rate and the congestion control algorithm shift \hat{G}_T towards lower load levels, which exhibit lower loss rates. This behavior is similar to the DVB-RCS2 normative load control, when G_T is forced to a value lower than G^* : TCP congestion control algorithm leads \hat{G}_T to oscillate around a stable equilibrium point without any centralized entity that adopts complex load control strategies.


We recall that $\lambda = G/N$ is the offered load per RCST. Figure 7b shows that, in the presence of few RCSTs (up to 100, approximately), a reduced number of larger time-slots (as for WF 14) should be preferred, in order to allow each RCST to take as much advantage as possible of the available resources. On the other hand, when the number of RCSTs increases, a WF that offers more time-slots (as for WF 3) should be preferred in order to allow a larger population to use the channel; in fact, even if λ shows little value variations among 100 and 400 RCSTs, the aggregate throughput in Figure 7a for the same load range is substantially larger or, at least, equal. It is worth noting in Figure 7b that λ has approximately the same value for $N > 400$, independently from WF and MSS. Finally, Figure 7b confirms that the scenario WF 14 and MSS = 23 bytes entails low performance and a large sub-utilization of available resources: thus, the choice of a WF and a MSS value which leads to $f > 1$ should be avoided.

TABLE VI: WF 14, MSS = 173 bytes, 64 time-slots per RA block. Relative error η when comparing the accuracy of the model presented in Section IV and, in the last two columns, when plugging the *BLR model* in (19) and (25)

#RCSTs	<i>NewRenoSAT</i> _{noTO}	<i>NewRenoSAT</i>	<i>BLRmodel</i> _{noTO}	<i>BLRmodel</i>
30	0.028	0.031	0.063	0.069
40	0.002	0.002	0.058	0.067
50	0.001	0.007	0.040	0.047
60	0.029	0.036	0.015	0.021
70	0.023	0.015	0.046	0.051



(a) Aggregate normalized MAC throughput and aggregate offered load vs number of RCSTs



(b) Offered load per RCST λ vs number of RCSTs

Fig. 7: MAC offered load and MAC throughput for different WFs and MSSs

VII. CONCLUSIONS

In this work, the performance of TCP NewReno has been analyzed over a random access satellite channel. Our proposed throughput estimation model, namely NewRenoSAT, accurately fits the simulation results for the satellite scenario under consideration, where losses are only due to collisions. Furthermore, our approach has been compared with other models in literature, showing that it achieves a lower estimation error. Our simulation results support the hypothesis that collisions can be considered as random independent events on a RA channel. Furthermore, a simple but effective model has been provided to estimate the loss event rate p and the segment loss rate q at transport layer; this is the *BLR model* that has shown a close match with simulation results, when used in the NewRenoSAT model.

In addition, the use of TCP enforces stability in RA channels, removing the need of the DVB-RCS2 load control mechanism, which is strongly scenario-dependent. TCP acts in a distributed way, without the need of a centralized entity in charge of enforcing potentially complex load control strategies. The use of different waveforms and MSS sizes has been analyzed, in order to identify the combination providing the best performance in the IoT/M2M scenario under consideration, also depending on the number of RCSTs. Furthermore, any system configurations that lead to the fragmentation of TCP

segments at MAC layer should be avoided because of the poor offered performance. In conclusions, the use of a TCP-based application protocol MQTT-like can support a large population in a satellite system, exploiting the advantages offered by the PUB/SUB paradigm in this context.

REFERENCES

- [1] J. Kim, J. Lee, J. Kim, and J. Yun, "M2M service platforms: survey, issues, and enabling technologies," *Communications Surveys & Tutorials, IEEE*, vol. 16, no. 1, pp. 61–76, 2014.
- [2] O. D. R. Herrero and R. De Gaudenzi, "High efficiency satellite multiple access scheme for machine-to-machine communications," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 48, no. 4, pp. 2961–2989, 2012.
- [3] S. Scalise, C. P. Niebla, R. De Gaudenzi, O. del Rio Herrero, D. Finocchi, and A. Arcidiacono, "S-MIM: a novel radio interface for efficient messaging services over satellite," *Communications Magazine, IEEE*, vol. 51, no. 3, pp. 119–125, 2013.
- [4] P. Cerwall, P. Jonsson, R. Möller, S. Bävertoft, S. Carson, I. Godor, P. Kersch, A. Källemark, G. Lemne, and P. Lindberg, "Ericsson mobility report," *Ericsson [Online]*. Available: www.ericsson.com/res/docs/2015/ericsson-mobility-report-june-2015.pdf, June 2015.
- [5] IDATE and UMTS, "Forum report 44," *Mobile Traffic Forecasts*, vol. 2020, 2010.
- [6] A. Laya, L. Alonso, and J. Alonso-Zarate, "Is the random access channel of LTE and LTE-A suitable for M2M communications? A survey of alternatives," *Communications Surveys & Tutorials, IEEE*, vol. 16, no. 1, pp. 4–16, 2014.
- [7] "Second generation, D.V.B. Part2: Lower layers for satellite standard," ETSI EN 301 545-2, 2012.

- [8] E. Casini, R. D. Gaudenzi, and O. Herrero, "Contention Resolution Diversity Slotted ALOHA (CRDSA): An enhanced random access scheme for satellite access packet networks," *IEEE Transactions on Wireless Communications*, vol. 6, no. 4, pp. 1408–1419, 2007.
- [9] O. del Rio Herrero and R. De Gaudenzi, "A high-performance MAC protocol for consumer broadband satellite systems," in *27th IET and AIAA International Communications Satellite Systems Conference (IC-SSC)*. IET, 2009, pp. 512–527.
- [10] M. Collina, M. Bartolucci, A. Vanelli-Coralli, and G. E. Corazza, "Internet of things application layer protocol analysis over error and delay prone links," in *Advanced Satellite Multimedia Systems Conference and the 13th Signal Processing for Space Communications Workshop (ASMS/SPSC)*, 2014 7th. IEEE, 2014, pp. 398–404.
- [11] J. Padhye, V. Firoiu, D. F. Towsley, and J. F. Kurose, "Modeling TCP Reno performance: a simple model and its empirical validation," *IEEE/ACM Transactions on Networking (ToN)*, vol. 8, no. 2, pp. 133–145, 2000.
- [12] N. Parvez, A. Mahanti, and C. Williamson, "An analytic throughput model for TCP NewReno," *IEEE/ACM Transactions on Networking (ToN)*, vol. 18, no. 2, pp. 448–461, 2010.
- [13] R. Dunaysev, K. Avrachenkov, Y. Koucheryavy, and J. Harju, "An Analytical Comparison of the Slow-but-Steady and Impatient Variants of TCP New Reno," in *Wired/Wireless Internet Communications*. Springer, 2007, pp. 30–42.
- [14] "The Network Simulator NS-3." [Online]. Available: <http://www.nsnam.org>
- [15] N. Cardwell, S. Savage, and T. Anderson, "Modeling TCP latency," in *IEEE INFOCOM*, 2000, pp. 1724–1751.
- [16] B. Sikdar, S. Kalyanaraman, and K. Vastola, "Analytic models for the latency and steady-state throughput of TCP Tahoe, Reno, and SACK," *IEEE/ACM Transactions on Networking*, vol. 11, no. 6, pp. 959–971, Dec 2003.
- [17] K. Zhou, K. Yeung, and V. Li, "Throughput modeling of TCP with slow-start and fast recovery," in *Global Telecommunications Conference, 2005. GLOBECOM '05. IEEE*, vol. 1, Nov 2005, p. 5 pp.
- [18] Z. Chen, T. Bu, M. Ammar, and D. Towsley, "Comments on modeling TCP Reno performance: a simple model and its empirical validation," *IEEE/ACM Transactions on Networking (ToN)*, vol. 14, no. 2, pp. 451–453, 2006.
- [19] C. Liu and E. Modiano, "An analysis of TCP over random access satellite links," in *Wireless Communications and Networking Conference, 2004. WCNC. 2004 IEEE*, vol. 4, March 2004, pp. 2033–2040 Vol.4.
- [20] G. L. Choudhury and S. S. Rappaport, "Diversity ALOHA - A Random Access Scheme for Satellite Communications," *Communications, IEEE Transactions on*, vol. 31, no. 3, pp. 450–457, 1983.
- [21] N. Celandroni and R. Secchi, "Suitability of DAMA and contention-based satellite access schemes for TCP traffic in mobile DVB-RCS," *IEEE Transactions on Vehicular Technology*, vol. 58, no. 4, pp. 1836–1845, May 2009.
- [22] N. Celandroni, F. Davoli, E. Ferro, and A. Gotta, "On Elastic Traffic via Contention Resolution Diversity Slotted Aloha Satellite Access," *International Journal of Communication Systems*, 2014.
- [23] M. Bacco, T. De Cola, G. Giambene, and A. Gotta, "Advances on Elastic Traffic via M2M Satellite User Terminals," in *2015 International Symposium on Wireless Communication Systems (ISWCS)*. IEEE, 2015, pp. 226–230.
- [24] M. Bacco, T. De Cola, and A. Gotta, "TCP New Reno over DVB-RCS2 Random Access Links: Performance Analysis and Throughput Estimation," in *2015 IEEE Global Communications Conference (GLOBECOM)*. IEEE, 2015, pp. 1–6.
- [25] M. De Sanctis, E. Cianca, I. Bisio, G. Araniti, and R. Prasad, "Satellite Communications Supporting Internet of Remote Things," *IEEE Internet of Things*, vol. 3, no. 1, pp. 113–123, 2016.
- [26] C. Joo and S. Bahk, "Start-up transition behaviour of TCP NewReno," *Electronics Letters*, vol. 35, no. 21, pp. 1818–1820, Oct 1999.
- [27] N. Cardwell, S. Savage, and T. Anderson, "Modeling TCP latency," in *IEEE INFOCOM, 19th Annual Joint Conference of the IEEE Computer and Communications Societies*, vol. 3. IEEE, 2000, pp. 1742–1751.
- [28] J. Puttonen, S. Rantanen, F. Laakso, J. Kurjenniemi, K. Aho, and G. Acar, "Satellite model for Network Simulator 3," in *Proceedings of the 7th International ICST Conference on Simulation Tools and Techniques*, Lisboa, Portugal, March 2014, pp. 86–91.
- [29] M. Bacco, P. Cassarà, E. Ferro, and A. Gotta, "Generalized Encoding CRDSA: Maximizing Throughput in Enhanced Random Access Schemes for Satellite," in *Personal Satellite Services*. Springer, 2013, pp. 115–122.
- [30] A. Meloni and M. Murrone, "Random Access in DVB-RCS2: Design and Dynamic Control for Congestion Avoidance," *IEEE Transactions on Broadcasting*, vol. 60, no. 1, pp. 16–28, 2014.
- [31] A. Munari, G. Acar, C. Kissling, M. Berio, and H. P. Lexow, "Multiple access in DVB-RCS2 user uplinks," *International Journal of Satellite Communications and Networking*, vol. 32, no. 5, pp. 359–376, 2014.